Photonic

WHITE PAPER

SHYPS to Shor's

A Call for Distributed QRE

DECEMBER 2025

Executive Summary

Quantum computing has moved from a question of *if* to a question of *when*. While there is no single answer as to *when* certain applications will be available, there are tools that provide insight. Quantum Resource Estimation (QRE) combines a variety of factors that impact quantum computation to create an estimate of the amount of quantum resources and the time needed to run a specific application on specific hardware.

Accurate QRE is important as a benchmark between different qubit modalities and system architectures. It is also a vital metric to track progress toward quantum utility – both *when* certain applications will become viable, and to measure the impact of improvements in hardware, software, and algorithms.

The issue is that most QRE analyses to date have assumed a monolithic architecture – a million or more qubits operating seamlessly in a single module for days at a time. Given the massive engineering challenge this represents, networking a series of modules to operate in unison may be a more practical and expedient path to implementing applications that require many thousands to millions of qubits. Increasingly, distributed quantum computing is recognized as a solution to the challenge of scale, but this approach can be expensive to implement if not integrated into the design from the start.

This white paper presents the first QRE analysis for Shor's algorithm to be precisely calculated for a **distributed** quantum computing architecture running a high-rate QLDPC error-correcting code. It forecasts a level of resources which is competitive with similar QRE analyses using the same algorithm approach, and it does so on a system that was created to deliver distributed quantum computing by design.

Photonic's Entanglement First™ Architecture

First Fully Costed Quantum Resource Estimation
with SHYPS QLDPC for Shor's algorithm:

- **7 million qubits**
- **3.9 days run time**

The need to include distributed costs in QRE analyses for comparative and roadmap purposes cannot be overlooked. Distributed quantum computing offers an alternative to monolithic systems by scaling systems linking optimally-sized, high-fidelity qubit processors via quantum interconnects. A distributed scaling paradigm overcomes the cross-talk, cryogenic, I/O, and fabrication constraints of monolithic architectures—and is likely to reach commercially relevant qubit counts much sooner. QRE that accounts for full system costs allows for more accurate assessments of each architecture's path to commercial value.

What is QRE, and Why is Distributed QRE Important?

Quantum computing holds the promise of transforming industries from chemistry, materials science, and drug discovery to cybersecurity, logistics, and finance. Delivering on that promise requires quantum computing resources capable of running commercially relevant quantum algorithms. While systems of this scale are not available yet, there is a way to plan for their arrival. Quantum Resource Estimation (QRE) provides – on a per-architecture basis – an estimate of the quantum resources required to run a particular algorithm.

Quantum Resource Estimation matters in understanding the speed and scale requirements for commercially relevant quantum computers.

Development of accurate QRE benchmarking methods is important for several reasons:

1. It allows comparisons of offerings from different quantum computer modalities, architectures, and algorithm implementations

2. It provides more accurate assessments of when key applications may become available based on hardware roadmaps

3. It creates a metric to gauge the impact of improvements in hardware, software, and algorithms

Types of QRE Tools

Not all QRE tools are made equal, though they typically report the same metrics: the **number of physical qubits** and **run time** required to run a specific algorithm on a given quantum computing system (number of logical qubits, and space-time complexity may also be reported). A full quantum computing stack involves complex, interdependent levels of hardware, software, and algorithms, making it challenging to compare the anticipated performance of different modalities or systems.

Looking at how QRE tools differ in input, and how accurate and comprehensive their estimations, it becomes clear that across all current approaches, one important consideration has been overlooked – **the inevitable cost of networking between modules, or distributed QRE**.

Logical QRE is the least accurate method of estimation. It assumes that all qubits are perfect (logical qubits) and does a very high-level estimate of how many 'perfect' qubits are required and how many operations the algorithm will require. Very often it only counts the most time-consuming type of gates (such as Toffoli or T gates) and ignores all other operations. It gives a very high-level estimate and does not take the architecture or compilation strategies into consideration, yet it is still useful for ranking various use cases for prioritization purposes. It is also the easiest method to use, making it useful as a first step in evaluating a large number of potential applications or use cases.

Physical QRE is a more advanced approach. It acknowledges that qubits are imperfect and require error correction. It factors in the use of particular error correcting codes to encode physical qubits into logical ones and does a proper estimation of those overheads. It also considers the fact that logical operations will need to be decomposed into basic operations on the physical qubits to arrive at a more realistic run time. The downside of physical QRE analysis is the additional upfront preparation work and algorithm optimization required – mapping the use case to a quantum circuit, careful selection of error correction code parameters, etc.

Hardware-Aware QRE is the most advanced and most resource-intensive form of estimation. In addition to accounting for concrete implementation of an algorithm and error correction parameters, this approach takes into consideration each system's particular architecture, qubit topology and performance, parallelization, realistic noise models, and other factors. This results in the most accurate estimation of required resources and run times.

The Need for Hardware-Aware, Distributed QRE

A key assumption made by most QRE tools is that all qubits will operate within a single module. Specifically, they assume only local interconnects and therefore do not account for the performance of interconnects between the modules – which cannot be assumed to operate the same as local interconnects in every architecture. Distributed QRE tools include performance of interconnects between the modules in their models.

The monolithic bias results in overly optimistic estimates for commercially or cryptographically important applications. This assumption may be reasonable for demonstration-scale circuits or early applications of quantum computing that can be run on smaller-scale systems, but the most impactful applications of quantum computing will require hundreds to thousands of logical qubits working in unison with extremely low logical error rates.

Factoring in error correction, this translates to millions of physical qubits. For most proposed quantum computing architectures, engineering a monolithic quantum computer of this size is a significantly more difficult path to scale. Constructing a network of modules is therefore the most cost-effective, and likely the fastest, path to commercial-scale computing. For scalable, reliable, high-performance systems, **distributed quantum computing is essential**.

As mentioned above, the vast number of logical qubits needed to operate advanced algorithms will require efficient error correction. Current physical-to-logical overheads for application-grade logical qubits using traditional surface codes are roughly 1000:1. However, the SYHPS family of Quantum Low-Density Parity Check (QLDPC) codes, constructed by Photonic and featuring an efficient logic implementation, can bring these overheads down to 100:1 or better by leveraging non- local connectivity between qubits – an improvement of 10x or more. Only architectures with long-range connectivity can take advantage of these codes. **Efficient, low overhead SHYPS QLDPC codes are distributed quantum computing compatible.**

Non-local connectivity enables the use of more efficient error correction codes that are compatible with distributed quantum computing.

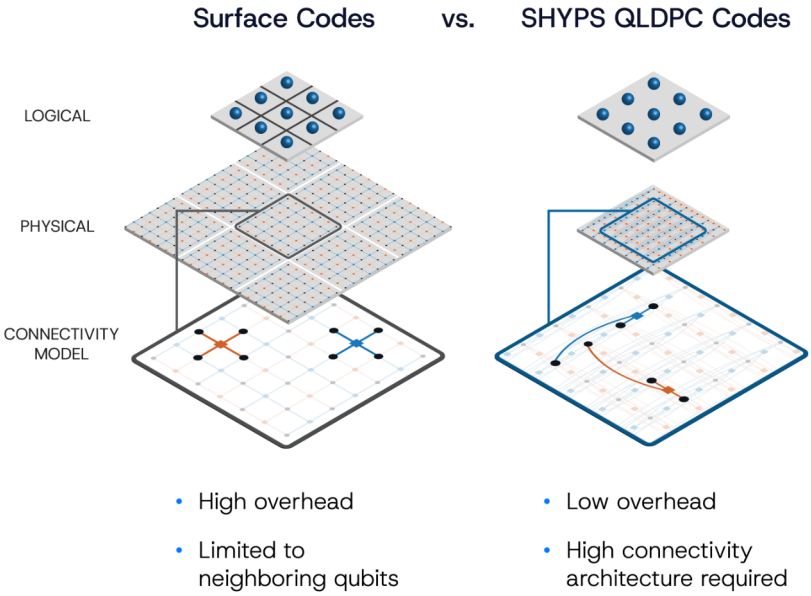

Given the inevitability and advantages of distributed quantum computing, why don't current QRE tools factor in the need for – and associated resource cost of – networking between qubits and modules? As quantum algorithms and error correction approaches continue to evolve, it is expected that resource requirements will fall dramatically over time. But these reductions will not be of the magnitude required to negate the need for inter-module connectivity. **Distributed quantum computing is needed for commercial scale.**

Photonic's Entanglement First™ Architecture

In the race to commercial deployment, the metrics have shifted – first from qubit count to logical qubit count. Increasingly, there is also a shift from single module to networkable systems for the achievement of commercially relevant scale. [GQI 2024] Large-scale, distributed, fault tolerant quantum computing is the foundation of Photonic's silicon colour centre based architecture, which was proposed in 2020 [2020 PRX Quantum paper] and described in detail shortly thereafter [2024 PRX Quantum paper].

Photonic's systems are based on an Entanglement First™ architecture designed for distributed quantum computing.

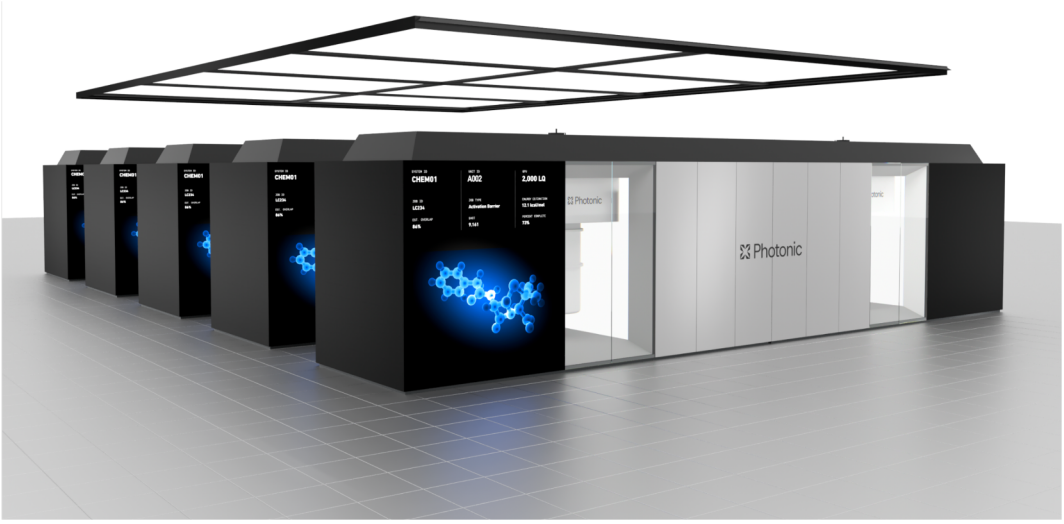

In Photonic's Entanglement First™ architecture, system capacity is scaled by adding high-qubit-count modules linked by fast quantum interconnects. The architecture prioritizes entanglement; separate silicon-based modules are connected through optical fibre networks, which both allows fast and efficient distribution of entanglement and large-scale interconnected networks of modules capable of performing utility-scale quantum computing.

Photonic is scaling capacity with high-qubit-count modules linked by fast quantum interconnects.

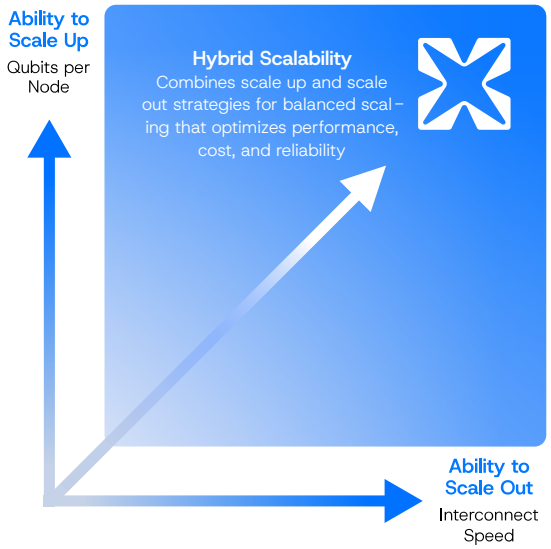

Shor's Algorithm: An Industry Benchmark

Shor's algorithm is frequently used as a benchmark algorithm due to its high relevance to the security of data and communications – as it will break common public key encryption standards when large scale quantum computers will become available. It is a well-defined computational task relevant to any architecture aiming for commercial scale, which makes it a good benchmark for evaluating the expected performance of a given quantum system architecture and its associated hardware parameters.

For several years, an implementation of Shor's algorithm devised by Gidney–Ekerå for superconducting qubits using surface code with lattice surgery has been considered best-in-class [Gidney–Ekerå 2021] and thus offers a good baseline for comparison of Photonic's architecture to other modalities. Here, we present the results of the first distributed QRE analysis of Shor's algorithm optimized for Photonic's Entanglement First™ architecture using SHYPS QLDPC codes, which inherently takes networking costs into account for a comprehensive estimate that is aligned with the final system implementation.

Optimizing Shor's Algorithm for QLDPC Codes

The dominant quantum cost in Shor's factoring algorithm comes from performing Quantum Phase Estimation implemented with many repeated quantum arithmetic subroutines. Building on the Gidney–Ekerå optimization techniques (originally developed for surface-code architectures) that minimize both the number of logical qubits and the number of expensive non-Clifford gates required to perform this arithmetic, we adapt and further optimize them for Photonic's Entanglement First™ architecture and distributed networked execution, targeting the factoring of 2048-bit RSA numbers.

A detailed calculation of the quantum resources required for this implementation of Shor's algorithm – including the overhead associated with multi-qubit operations, error-corrected non-Clifford gates, and the auxiliary qubits needed to support them – was completed using Photonic's hardware performance factors and leveraging the block nature of the proprietary Subsystem Hypergraph Product Simplex (SHYPS) QLDPC codes to enable time-optimal quantum computation.

Photonic's distributed QRE results are shown below relative to the Gidney–Ekerå 2021 benchmark [Gidney–Ekerå 2021]. It is important to note that the Gidney–Ekerå resource estimation was performed on surface codes, whereas Photonic's estimates are optimized for QLDPC codes. As a result, the inherent high connectivity of Photonic's architecture dramatically reduces the routing overhead historically assumed for modalities using surface codes.

Shor's Algorithm

Photonic Distributed QRE with QLDPC on optically linked silicon spin qubits:

7 million qubits **3.9 days run time**

Gidney–Ekerå Monolithic QRE with Surface Code on superconducting qubits:

20 million qubits **8 hours run time**

Using a networked approach and SHYPS QLDPC codes, Photonic's system could feasibly factor 2048-bit numbers using **7 million qubits in under 4 days**. Not only is this resource-competitive with the other approaches shown, but this estimate naturally incorporates distributed network costs rather than assuming a monolithic quantum computer of 19–20 million qubits. Furthermore, these results represent the first precise QRE calculation for Shor's algorithm on a high-rate QLDPC error-correcting code (all other precise calculations have been done on surface codes, and other QLDPC calculations have utilized more estimation).

It is worth noting that an alternate implementation of Shor's algorithm has been proposed recently for superconducting qubits to reduce QRE to 1 million qubits, with a run time of 7 days [Gidney, 2025]. However, this QRE analysis also assumes a monolithic architecture, and as a result, excludes distributed costs. So, while it involves fewer physical qubits, it requires them to run longer, all operating in a single module. In contrast, Photonic's initial implementation of QLDPC codes for Shor's algorithm automatically incorporates distributed network costs. It has not yet been fully optimized to realize the full benefits of recent advances and therefore has significant opportunities for further resource reductions and speedups through additional compiler optimization, algorithm improvements, and improved QLDPC codes with even better logical error rates.

Taken together, these results clearly demonstrate the importance of comparing QRE analyses for different quantum computing systems, understanding the assumptions used, and the importance of determining whether networking costs have been considered.

Distributed QRE Maps a More Realistic Future

Quantum Resource Estimation provides an estimate of the quantum resources required to run a particular algorithm on a given architecture. This process helps identify which use cases will become practical in the near term, what kind of quantum computer capabilities will be required to address each application, and track progress towards them.

In doing so, the need to include distributed costs for comparative and roadmap purposes cannot be overlooked. The path to commercial value quantum can be accelerated with a networked approach to quantum computing, and distributed QRE provides the means to demonstrate this advantage.

While existing QRE analyses take into account many factors, networking costs are not typically included, making these estimates dependent on whether the required number of qubits can operate for the required run time in a single module without failure.

Photonic's architecture, in contrast, is distributed by design, and will allow multiple modules to work in concert to run a single algorithm. To reflect this capability, Photonic completed an initial QRE analysis for Shor's algorithm to be precisely calculated on a high-rate QLDPC error-correcting code, projecting a run time of 3.9 days using 7 million physical qubits. This projection is competitive with QRE analyses for other modalities that are also based on the same approach [Gidney–Ekerå 2021].

Future implementations leveraging recent algorithmic improvements [Gidney, 2025] will undoubtedly show even more competitive outcomes. At every step, distributed QRE will be more realistic for large scale applications, having factored in the need for distributed quantum computing.

References

Scalable Quantum Hardware, GQI Outlook Report (2024), *Global Quantum Intelligence*.

Bergeron, L., Chartrand, C., Kurkjian, A. T. K., Morse, K. J., Riemann, H., Abrosimov, N. V., ... & Simmons, S. (2020). Silicon-integrated telecommunications photon-spin interface. *PRX Quantum*, 1(2), 020301.

Simmons, S. (2024). Scalable fault-tolerant quantum technologies with silicon color centers. *PRX Quantum*, 5(1), 010102.

Gidney, C., & Ekerå, M. (2021). How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits. *Quantum*, 5, 433.

Gidney, C. (2025). How to factor 2048 bit RSA integers with less than a million noisy qubits. arXiv: 2520.15917

For more details on Photonic's approach to this world-changing technology, please visit www.photonic.com